Reviewer 1: Thanks for the feedback and for the suggestions as to how to make the paper clearer and the examples less intimidating. We'll work on that for the final version. • Re *"...how decomposing the polytope now allows it to be mapped?"* We are not sure we understand the question. If you meant "how does the decomposition help map the problem of computing an optimal correlated equilibrium onto an LP?", here is the answer. If the goal is to write an LP for computing an optimal correlated equilibrium in a triangle-free game, our crucial result is that $\Xi = \mathcal{V}$, and the scaled-extension decomposition can just be regarded as a detail in the proof of that equality. The LP for computing an optimal correlated equilibrium is of the form $\arg\max_{\boldsymbol{\xi}}\{\boldsymbol{c}^\top\boldsymbol{\xi} : \boldsymbol{A}\boldsymbol{\xi} + \boldsymbol{B}\boldsymbol{y} = \boldsymbol{0}, \boldsymbol{y} \geq \boldsymbol{0}, \boldsymbol{\xi} \in \Xi\}$, so once it is known that $\Xi = \mathcal{V}$, one can substitute the constraint $\boldsymbol{\xi} \in \Xi$ with $\boldsymbol{\xi} \in \mathcal{V}$, which corresponds to the (polynomially small) set of linear constraints given in Equation 1. • Re *"I wasn't sure what I was supposed to take away from the experiments"* As mentioned, the scaled-extension-based decomposition of $\mathcal{V}$ is not directly used to write a linear program. However, it can be used to construct an efficient regret minimization algorithm for the set $\mathcal{V}$, which in turn can be used to compute an EFCE, EFCCE, NFCCE [12] (see also Reviewer 5). In the experiments, we show that (i) we implemented and tested the decomposition algorithm, and that it is able to scale to large games, and (ii) that we were able to experimentally confirm that the regret minimization algorithm is more scalable than the linear programming approach in large games—both in terms of run time and memory—(using the same experimental setup as [12]).

Reviewer 2: Thanks for the constructive feedback and suggestions as to how to improve the presentation and make the paper less intimidating. We'll take all of them into account. • Re *"broader impact"* Thanks for the feedback, we agree with all your points. As you correctly recognized, we use the term "social welfare" to mean the sum of utilities of the players as is typical in the game theory literature, but as you rightly point out, that need not necessarily coincide with the societal notion of optimality. • Re *"scale of infeasibility"* The constraints in Equation (1) are not the only constraints in the LP. Also, the sum of entries of a generic vector that satisfies the constraints in (1) will sum to more than 1 usually (i.e., correlation plans have more than unit mass). Gurobi's initial point when using its implementation of the barrier algorithm might be way off in terms of satisfying the constraint of the (primal) LP, and that is why Gurobi reported that the first iterate had a very high infeasibility (defined in lines 754-761). The maximum payoff is 15. • Re *"free LP solver"* Thanks for the feedback. Gurobi is freely available for academic use, but we'll also mention the open-source (and slower, less numerically stable) alternative GLPK, unless the reviewers have better suggestions. • Re *"Goofspiel"* We are not aware of it having been solved for EFCE before; we'll check. We are definitely the first to compute *optimal* EFCE in it. We will point these out as additional minor contributions of our paper. We will add results for all values of k too, including plots with iteration count on the x-axis as suggested.

Reviewer 3: Thanks for your feedback. We strongly disagree that "*this paper just tells us that the work in Farina et al. [12] is straightforward to extend to the triangle-free game*" First we disagree that the extension is "straightforward"—and so seem to disagree the other reviewers too. Isolating and introducing triangle-freeness as a meaningful condition with consequences on the correlation structure of games is an important insight in itself. Extending the construction by Farina et al. to handle the more general structure is another (not easy) endeavor. Understanding how the existence of a decomposition relates to the integrality of the vertices of the von Stengel-Forges polytope is a third, separate result. On an independent note, we argue that whether the technical insights in a paper are simple or not in hindsight is not in itself a good metric for measuring the quality and potential impact of research. The other reviewers seem to think that the paper borders on too heavy technical contributions. You seem to suggest that it is too light on technical contribution. We strongly disagree with that. We hope the other reviews and our responses will convince you of the merits of the paper. • Re *"line 244"*: we meant leftmost. We'll fix the typo. • Re *"infeasibility"*. Thanks for the feedback. We'll make sure to define it in the body. Currently, a definition is available in the appendix, lines 754-761.

Reviewer 5: Thanks for pointing out typos and for the improvement suggestions! We'll apply them in the final version. • Re *"the decomposition should serve to find correlated equilibria more easily"* The equality $\mathcal{V} = \Xi$ that we proved via the decomposition makes it *possible at all* (not just *easier*) to *optimize* over the set of correlated equilibria in polynomial time. This is because $\mathcal{V}$ can be expressed as the intersection of few linear constraints, while Ξ in general cannot. (See also Reviewer 1.) • Re *"does not take full advantage of your approach in this work"* While our focus is on offering the first algorithms to compute an optimal EFCE (in triangle-free games), we agree that it is likely that they would also be faster than the prior methods for finding a feasible equilibrium: 1) Dudik&Gordon requires convergences of MCMC at each iteration, likely making it slow and numerically unstable; 2) Huang is believed to be of theoretical interest rather than practically fast—because it uses the ellipsoid method. However, we leave investigating that direction as future work. In the last bit of the experimental section, we were simply interested in recreating the same setup as [12] to show that their conclusion hold in triangle-free games as well. Finally, we note that D&G and Huang cannot be used to compute optimal equilibria. • Re *"the regret-minimization algorithm could not handle an objective?"* While we do not do that (and neither do the authors in [12]), the regret minimization method in [12] can find an equilibrium with a given lower bound on an objective. Hence, to optimize an objective one could, in theory, perform a binary search on the optimum objective value by running the regret minimization method several times with different lower bounds. However, that has never been tried in the literature. • Re *"...any tests...EFCE with an objective function..."* We do use an objective to compute the sets in Figure 3 (right). In a nutshell, what we did was to take a game and try objective functions of the form $\alpha \cdot$ payoff player $1 + \beta \cdot$ payoff player 2 (with $\|(\alpha, \beta)\|_2 = 1$). Given a choice of (α, β), we computed the optimal objective value $v_{\alpha,\beta}$ of any EFCE, EFCCE, NFCCE of the game by solving the corresponding linear program with Gurobi. This shows that the set of reachable payoffs must satisfy the inequality $\alpha \cdot$ payoff player $1 + \beta \cdot$ payoff player $2 \leq v_{\alpha,\beta}$. Taking the intersection of all these constraints yields the polytope of reachable payoffs of the game. We'll include these and more details in the final version. • Re *"did you observe any drop in performance...without an objective function?"* We did not extensively compare running the linear program with and without an objective on the same game instances, but our best guess informed from past experience with these equilibria is that the difference would be minimal when using Gurobi.